PAUL SCHERRER INSTITUT

PSI

Martin Brückner :: SLS Detectors Group :: Paul Scherrer Institut

An FPGA based high speed network processor –
Preprocessing detector images relieving data backend systems

21st IEEE Real Time Conference

PAUL SCHERRER INSTITUT

Outline

- Eiger and Jungfrau Detectors
- Receiving detector data
- FPGA based network processor

Detector Introduction

	Eiger	Jungfrau
	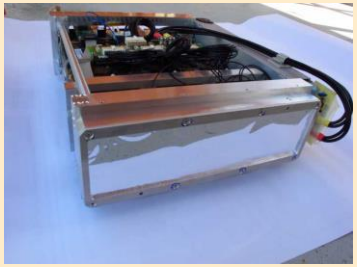	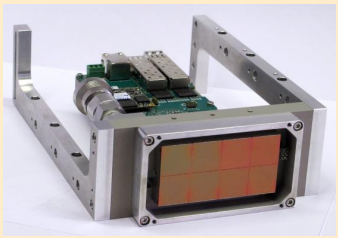
Type	Single photon counting (Increment pixel counter when photon hits pixel)	Charge integrating (Collect charge over exposure time)
Pixel size	75 x 75 µm²	75 x 75 µm²
Dynamic range	4bit/8bit/12bit/(32 bit)	3 gain stages
Max. frame rate	Burst: 22/11/6 kHz Continuous: 10/5/2.5/1.2 kHz	2.4 kHz continuous
Max. data rate per module	2x 10 Gbit/s (internal: 2x 23 Gbit/s)	2x 10 Gbit/s
Pixel per module	512*1024 Pixel	512*1024 Pixel

Detector Introduction

Eiger 9M	Jungfrau 10M	Jungfrau 16M
9 MPixel 18 modules 360 Gbit/s Running @SLS	10 MPixel 20 modules 336 Gbit/s Planned for SLS	16 MPixel 32 modules 27.2 Gbit/s In construction for SwissFEL

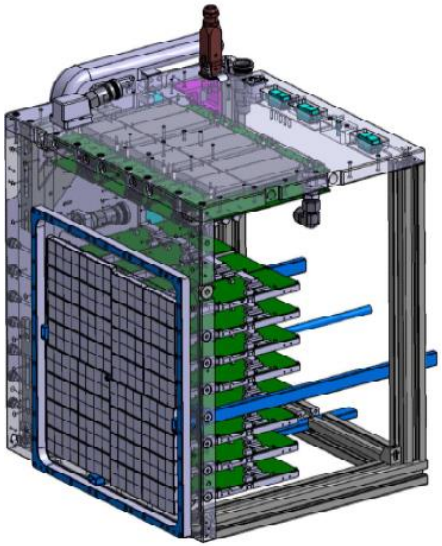

16M Jungfrau detector for SwissFEL

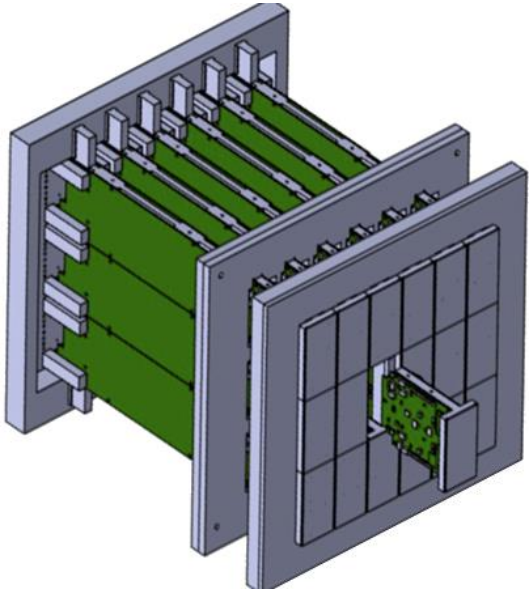

9M Eiger detector @cSAXS

Detector Introduction

- Module based detectors
  - Independent data taking
  - Parallel data sending
  - Larger detectors with more modules:
    - More pixels / Higher data rates
    - Same frame rate

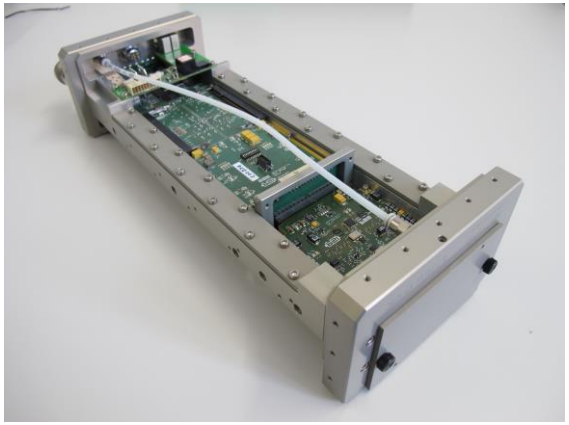

Eiger module

  - Sending UDP packets including
    - Frame Number
    - Packet Number
    - Module ID

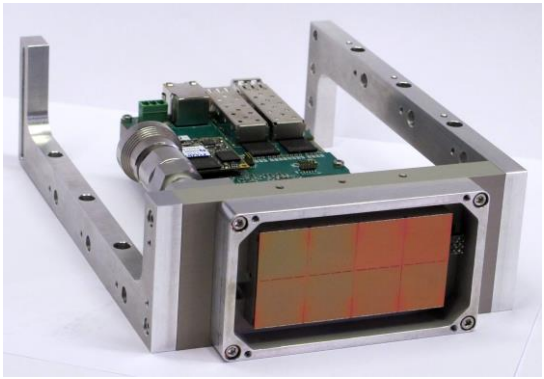

Jungfrau module

Receiving detector data

- Host receiver tasks
  - Receiving packets (in random order)
  - Reassembling the images using *Frame Number*, *Packet Number* and *Module ID* from the packet header
  - Keep track of missing packets (pad image)

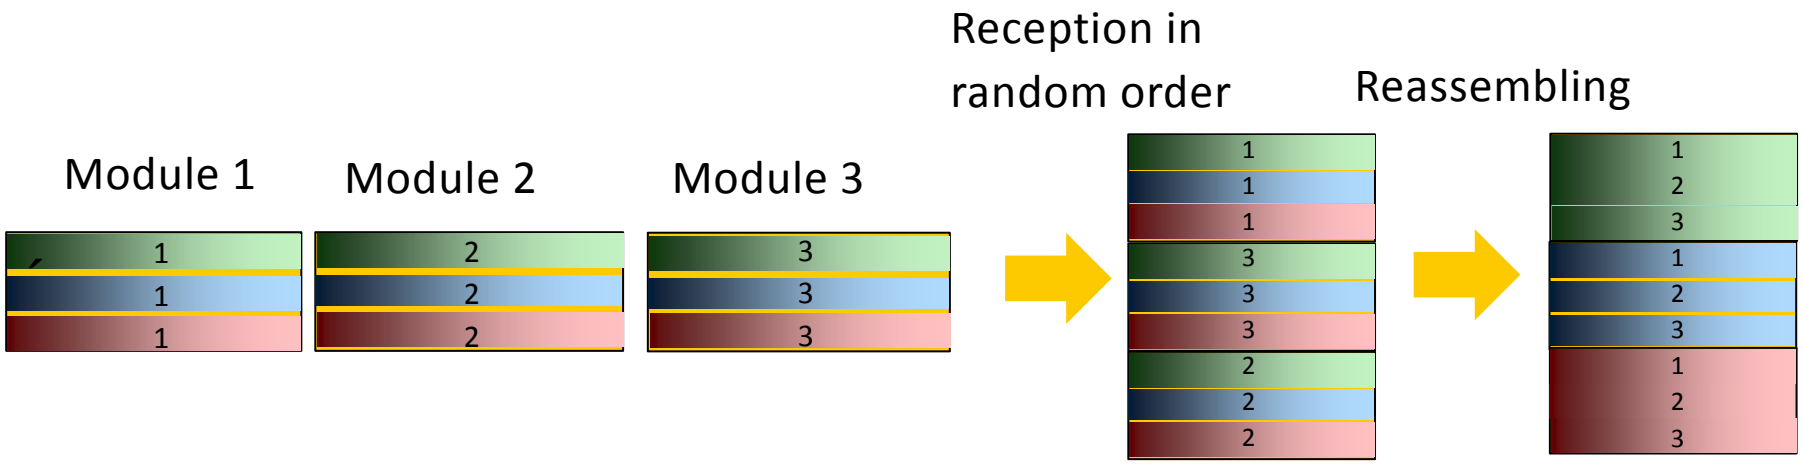

Eiger 4 Bit/Pixel	Eiger 8 Bit/Pixel	Eiger 16 Bit/Pixel	Eiger 32 Bit/Pixel
16 lines/packet	8 lines/packet	4 lines/packet	2 lines/packet

Receiving detector data

- For multi module systems performance tweaks have to be done
  - Bind receivers to specific CPU cores
  - Tune network buffer parameters
  - Use multiple network cards (one per CPU socket)

- For a single module Eiger detector at 16bit/2kHz on Dual Xeon E5646 @2.4GHz[1]
  - Receiver listeners: 4 cores with 80% load each[1]
  - 10Gb/s packet dispatch: 1 core with 70% load[1]

- Eiger 9M at the PSI cSAXS beamline (no high frame rate, yet)
  - 2 CPU cores per Eiger module → so 36 cores needed
  - 1 additional CPU core for interrupts

- Otherwise packet loss occurs

[1] Measured by Alejandro Homs (ESRF)

Receiving detector data

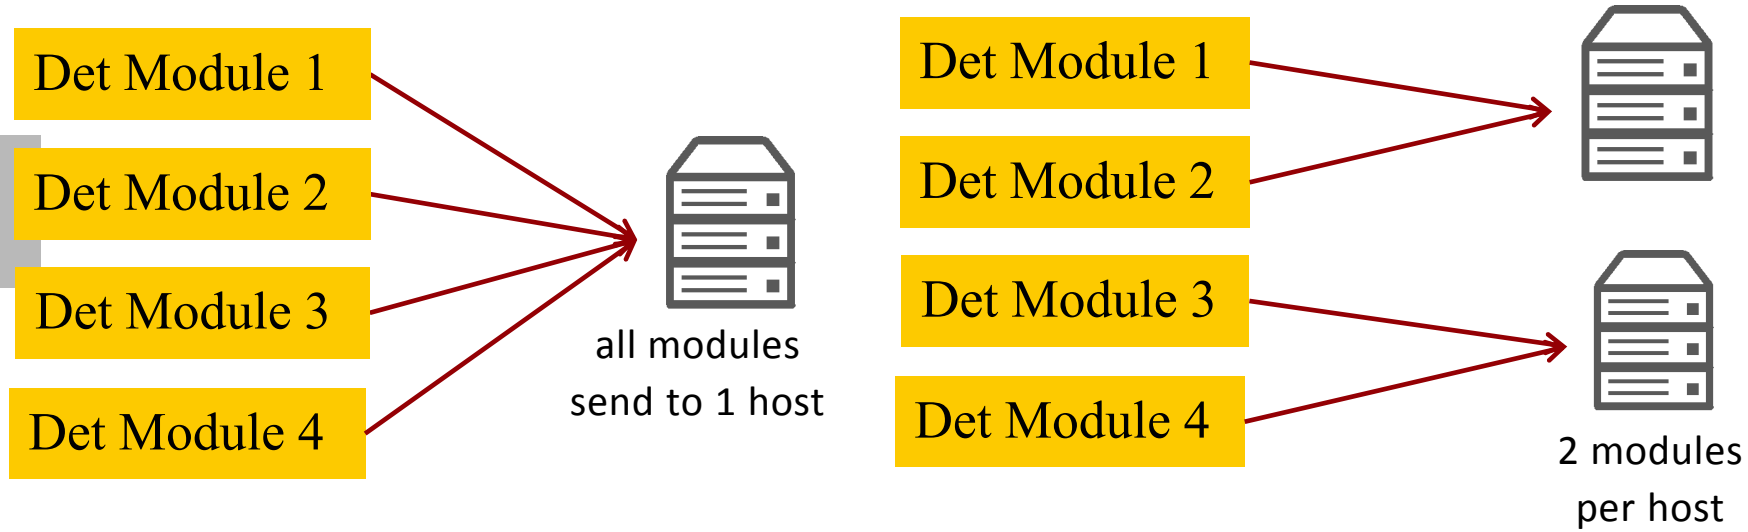

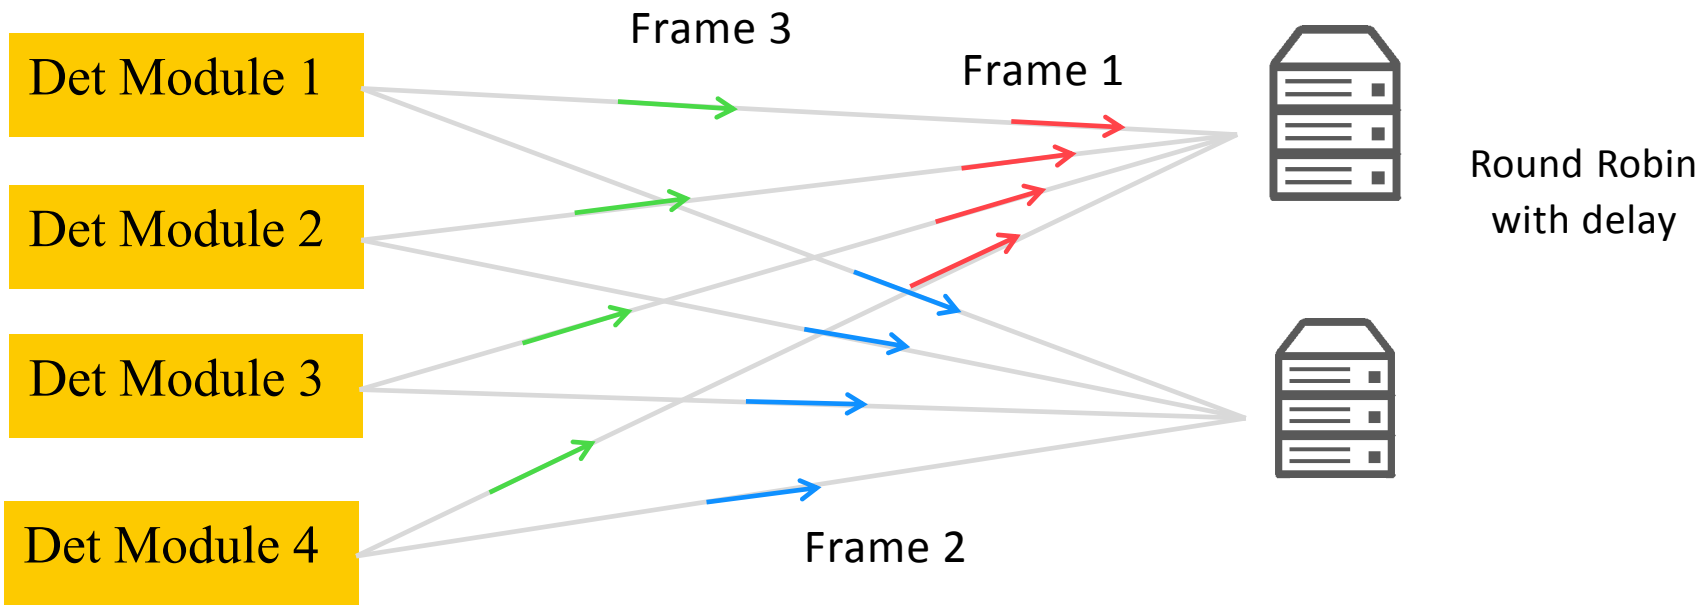

FPGA Based Network Processor

- FPGA based network processor
  - Can analyze packet headers and assemble images in host memory
  - Writes a proper image header for each image and discards packet headers
  - CPU receives a pointer to final image(s)

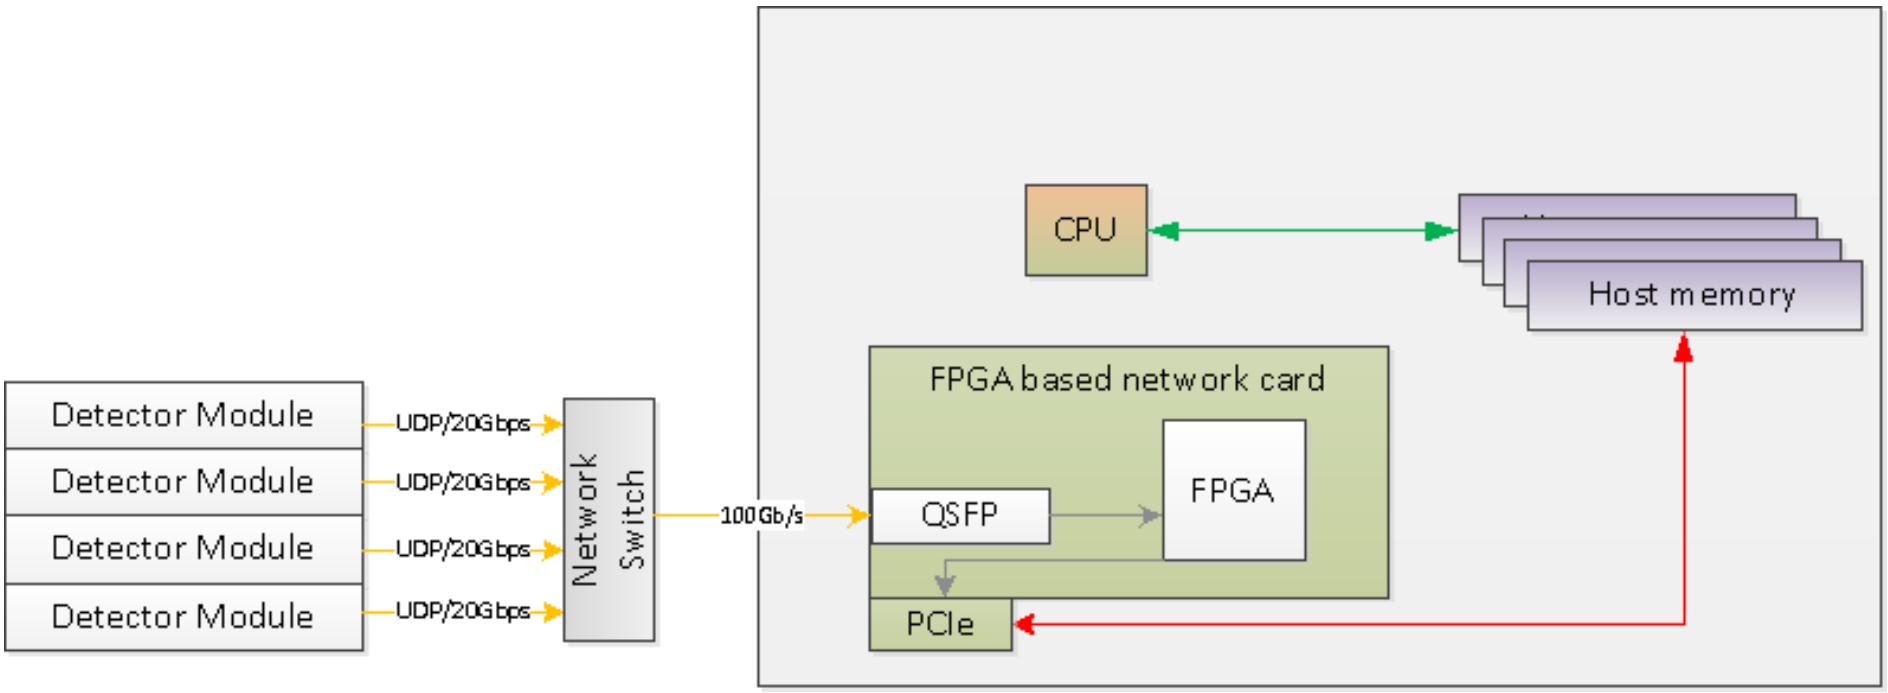

FPGA Based Network Processor – VCU108

Xilinx VCU108 Evaluation Board

- Virtex Ultrascale Hard IP cores
  - 100Gbit/s Ethernet
  - PCIe 3.0 x8 (64Gbit/s only)
  - No RS-FEC with 100Gbit/s
  - Cheap

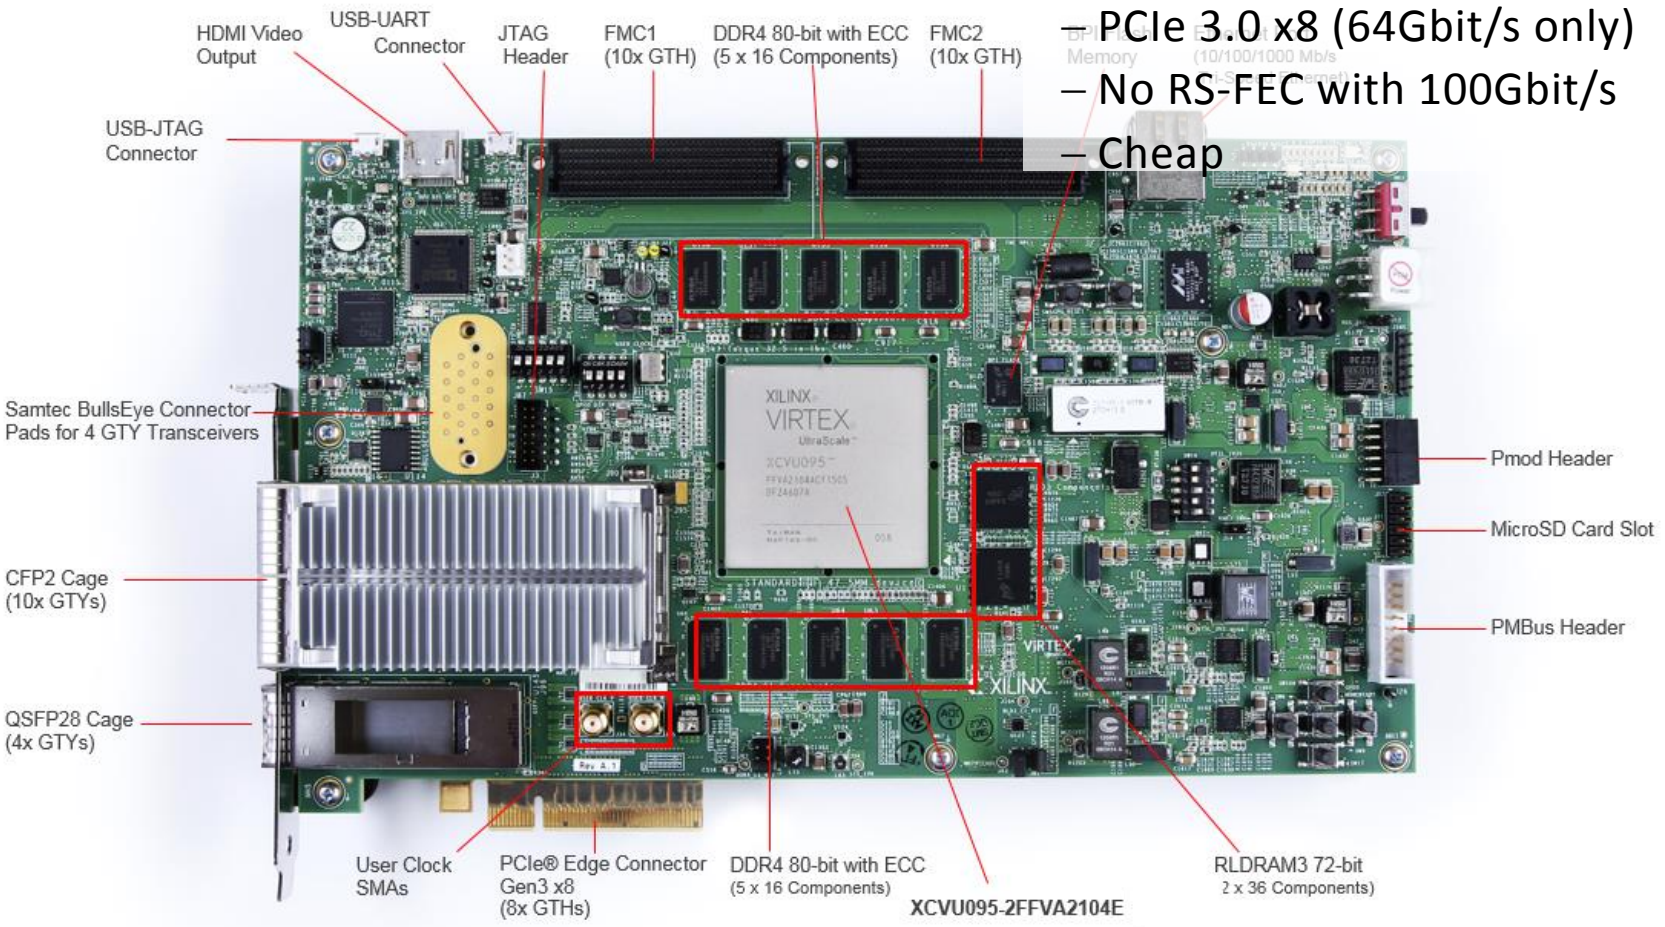

FPGA Based Network Processor – VCU1525

Xilinx VCU1525

- Virtex Ultrascale+ Hard IP cores
  - 2x 100Gbit/s Ethernet QSFP28
  - PCIe 3.0 x16 or PCIe 4.0 x8
  - RS-FEC included

Planned to buy

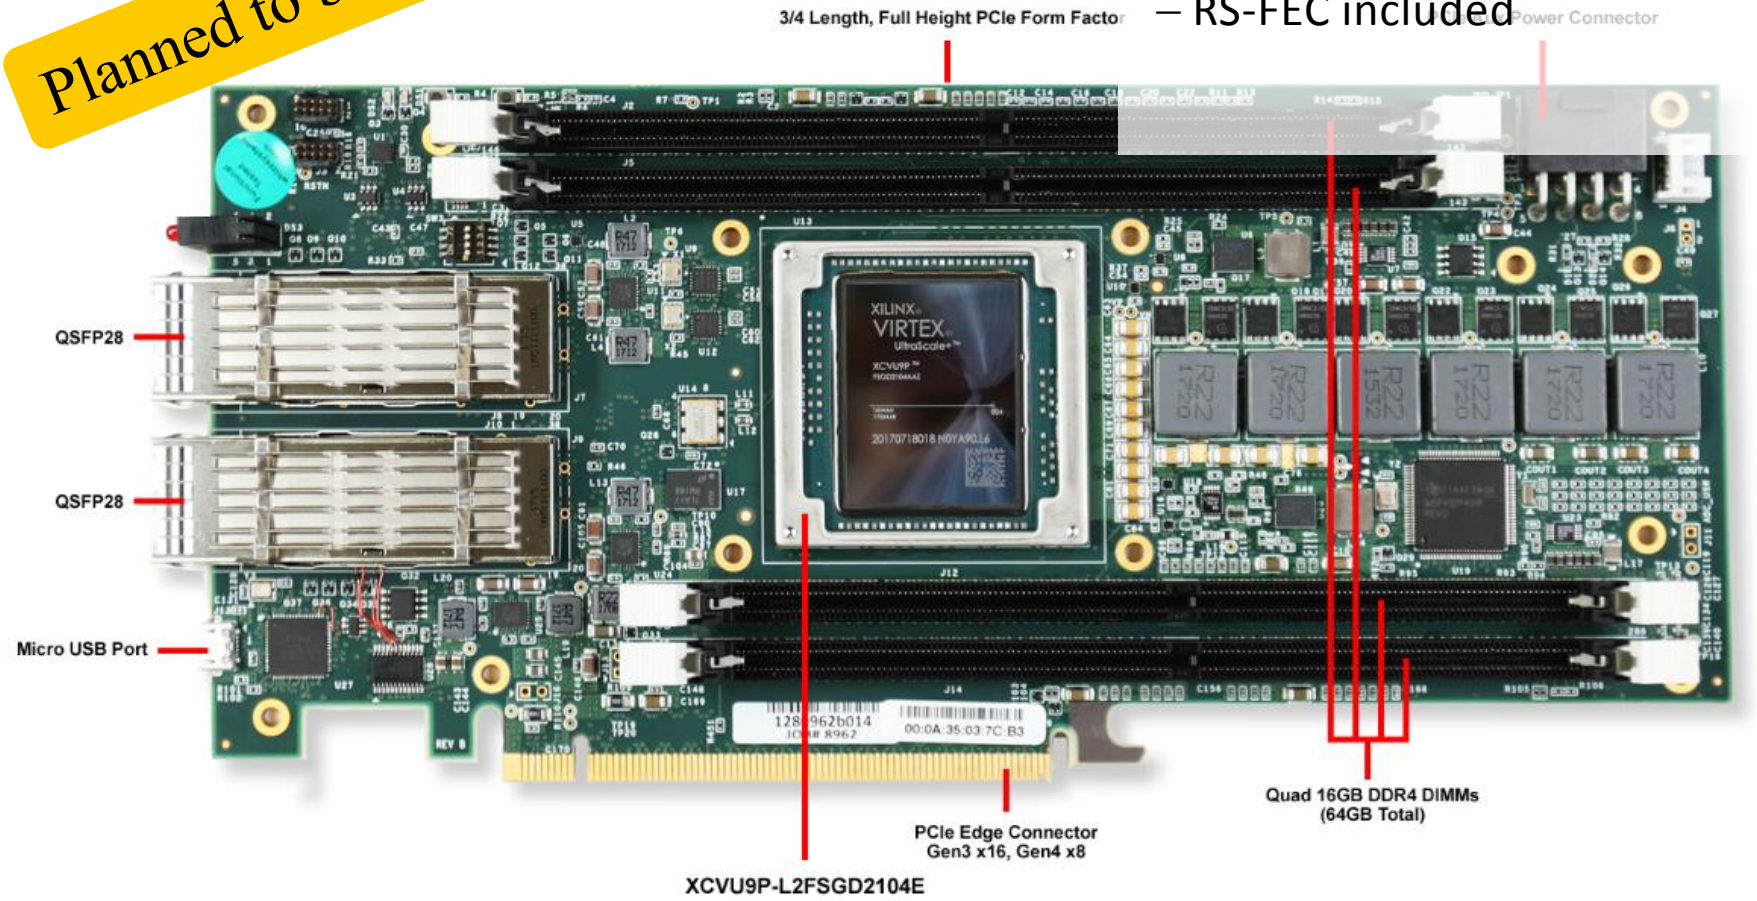

FPGA Design

- Converts the segmented LBus into 4 independent AXIS streams

- Generates memory addresses out of packet header information
- Can generate multiple addresses per packet (solves several lines/packet problem)

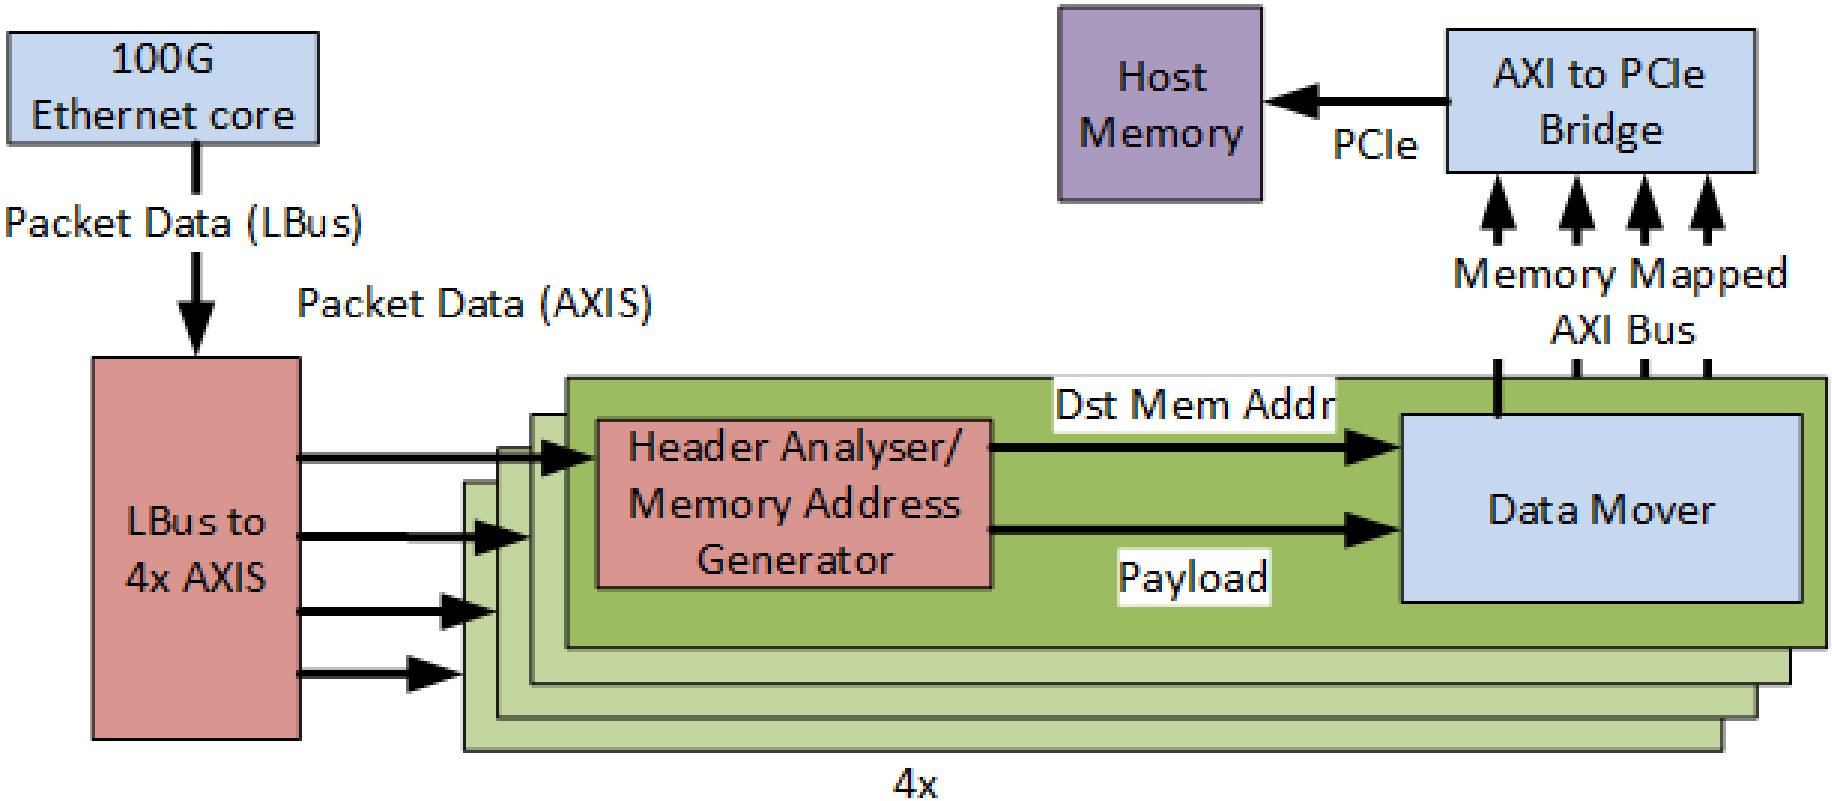

FPGA Based Network Processor – Testsetup

Test setup:

- Loopback SFP module
- Network Packet Generator sends Ethernet packets with 8192 Byte/packet

Result (obvious):

- 100G Ethernet works
- PCIe 3 x8 is the bottleneck

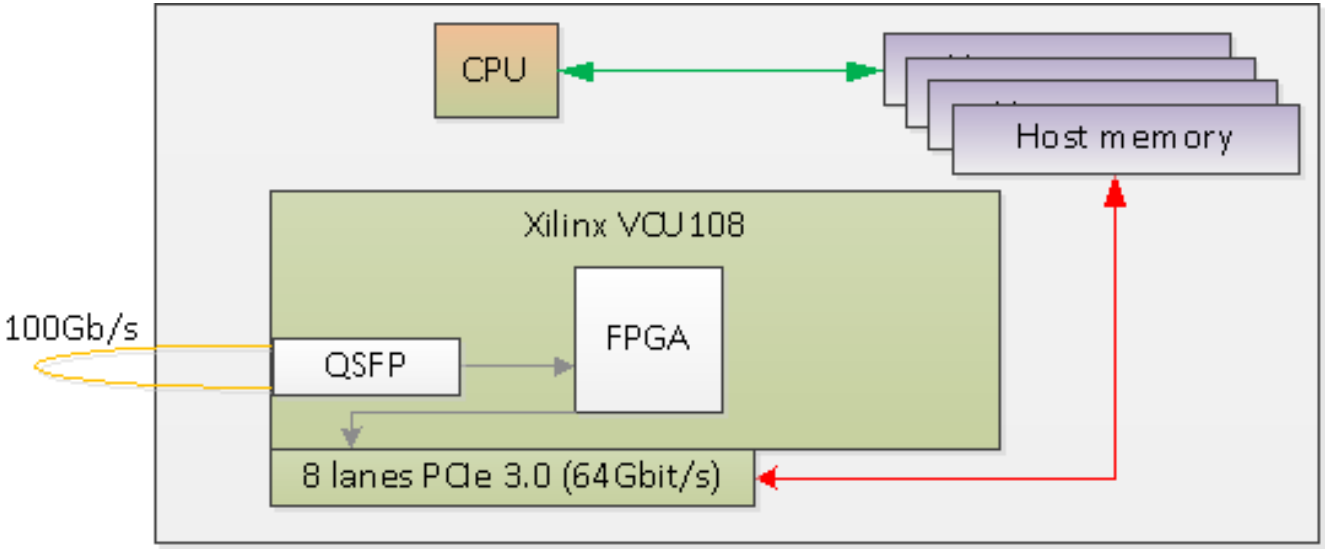

FPGA Based Network Processor – Testsetup

Test setup:

- Remove 100G core
- Generate packets in the FPGA

- Maximum: 64 Gbit/s PCIe 3.0 x8
  7.8 Gbyte/s (128/130 Encoding)

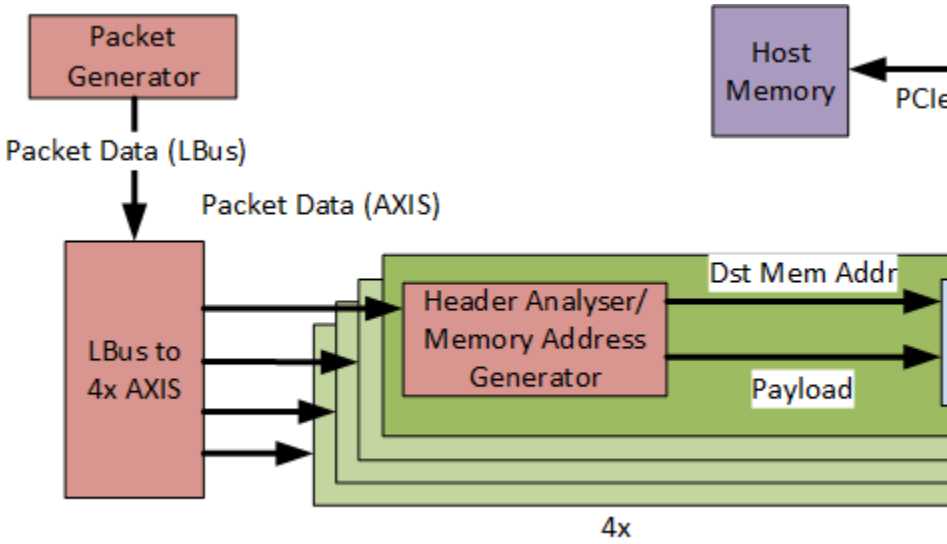

Results:

- No CPU load while FPGA copies data to host memory

Packet Size	4096 Byte	8192 Byte	16384 Byte
Speed	5.63 Gbyte/s	6.1 Gbyte/s	6.16 Gbyte/s

Measurement tool: Intel Performance Counter Monitor
https://github.com/opcm/pcm

Combine FPGA and GPU

- Direct data transfer FPGA → GPU using GPUDirect RDMA
- GPU can do
  - Charge-Sharing into photons conversion for Jungfrau
  - Cluster finding
  - Data compression
  - …

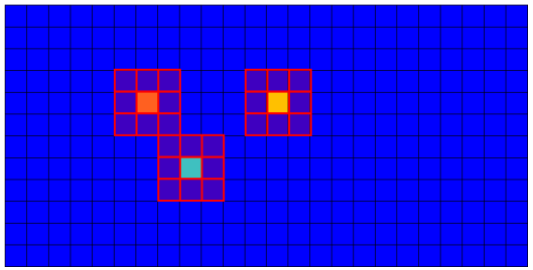

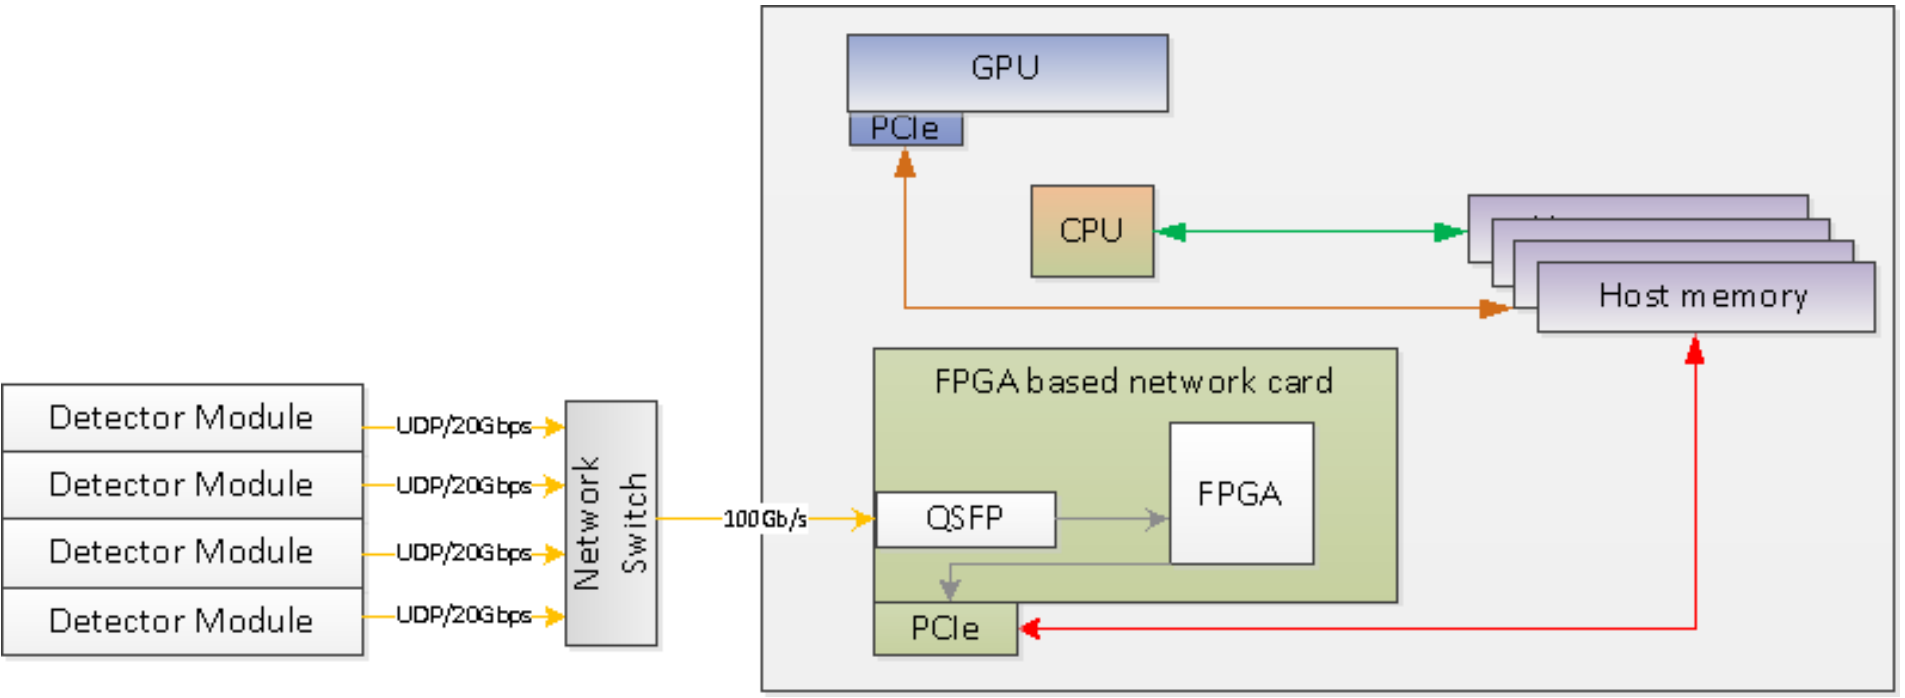

Conclusion

- Eiger and Jungfrau generate high data rates
- Difficult to receive them on normal PCs (high CPU load)
- FPGA can relieve the CPU and reassemble images in host memory
- Max data rate is around 80% of the maximum PCIe3 x8

- Future:
  - Try the VCU1525 with PCIe 3 x16
  - Combine FPGA receiving with GPU image processing

Thank You

The PSI SLS Detectors Group 2017